

HP Smart Array 641/642 Controller User Guide



July 2003 (Second Edition)
Part Number 309311-002

© 2003 Hewlett-Packard Development Company, L.P.

Microsoft®, Windows®, and Windows NT® are US registered trademarks of Microsoft Corporation.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided “as is” without warranty of any kind and is subject to change without notice. The warranties for HP products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

HP Smart Array 641/642 Controller User Guide

July 2003 (Second Edition)
Part Number 309311-002

Contents

About This Guide

Audience Assumptions	vii
Important Safety Information	vii
Symbols on Equipment	vii
Symbols in Text.....	viii
Related Documents.....	ix
Getting Help	ix
Technical Support	ix
HP Website	x
Authorized Reseller	x
Reader's Comments	x

Chapter 1

Installation Overview

Never-Used, Autoconfigurable Server	1-1
Never-Used, Non-Autoconfigurable Server	1-2
Previously Configured Server	1-3

Chapter 2

Installing the Hardware

Preparing the Server	2-1
Installing the Controller Board.....	2-2
Connecting Storage Devices.....	2-2
Connecting Internal Storage	2-3
Connecting External Storage (Model 642 only)	2-4
SCSI Cable Part Numbers.....	2-4

Chapter 3

Updating the Firmware

Chapter 4

Configuring the Server

Using RBSU	4-1
Using SCU	4-2

Chapter 5

Configuring an Array

Using ACU	5-3
Using ORCA.....	5-3
Configuration Procedure	5-4

Chapter 6

Installing the Device Drivers and Management Agents

Device Drivers	6-1
Management Agents	6-1

Chapter 7

Upgrading or Replacing the Cache

Appendix A

Regulatory Compliance Notices

Federal Communications Commission Notice	A-1
Class A Equipment.....	A-2
Class B Equipment	A-2
Declaration of Conformity for Products Marked with the FCC Logo, United States Only	A-3
Modifications	A-3
Cables	A-3
Canadian Notice (Avis Canadien)	A-4
Class A Equipment.....	A-4
Class B Equipment	A-4
European Union Notice	A-4
BSMI Notice.....	A-5
Japanese Notice.....	A-5

Korean Notices	A-5
Battery Replacement Notice	A-6

Appendix B

Electrostatic Discharge

Appendix C

Controller Specifications

Appendix D

Drive Arrays and Fault Tolerance

What Is a Drive Array?	D-1
Fault-Tolerance Methods.....	D-5
Hardware-Based Fault-Tolerance Methods	D-5
Alternative Fault-Tolerance Methods	D-11

Appendix E

Replacing, Moving, or Adding Hard Drives

Identifying the Status of a Hard Drive	E-1
Recognizing Hard Drive Failure	E-3
Effects of a Hard Drive Failure.....	E-3
Compromised Fault Tolerance.....	E-4
Recovering from Compromised Fault Tolerance.....	E-4
Replacing Hard Drives	E-5
Factors to Consider Before Replacing Hard Drives.....	E-5
Automatic Data Recovery (Rebuild)	E-7
Upgrading Hard Drive Capacity	E-8
Moving Drives and Arrays	E-9
Adding Drives	E-11

Appendix F

Probability of Logical Drive Failure

Appendix G

Troubleshooting

Index

About This Guide

This guide provides step-by-step instructions for installation, and reference information for troubleshooting, for the HP Smart Array 641 and 642 Controllers.

Audience Assumptions

This guide is for the person who installs, administers, and troubleshoots servers. HP assumes you are qualified in the servicing of computer equipment and trained in recognizing hazards in products with hazardous energy levels.

Important Safety Information

Before installing this product, read the *Important Safety Information* document included with the server.

Symbols on Equipment

The following symbols may be placed on equipment to indicate the presence of potentially hazardous conditions:



WARNING: This symbol, in conjunction with any of the following symbols, indicates the presence of a potential hazard. The potential for injury exists if warnings are not observed. Consult your documentation for specific details.



This symbol indicates the presence of hazardous energy circuits or electric shock hazards. Refer all servicing to qualified personnel.

WARNING: To reduce the risk of injury from electric shock hazards, do not open this enclosure. Refer all maintenance, upgrades, and servicing to qualified personnel.



This symbol indicates the presence of electric shock hazards. The area contains no user or field serviceable parts. Do not open for any reason.

WARNING: To reduce the risk of injury from electric shock hazards, do not open this enclosure.



This symbol indicates the presence of a hot surface or hot component. If this surface is contacted, the potential for injury exists.

WARNING: To reduce the risk of injury from a hot component, allow the surface to cool before touching it.

Symbols in Text

These symbols may be found in the text of this guide. They have the following meanings.



WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or loss of life.



CAUTION: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or loss of information.

IMPORTANT: Text set off in this manner presents essential information to explain a concept or complete a task.

NOTE: Text set off in this manner presents additional information to emphasize or supplement important points of the main text.

Related Documents

For additional information on the topics covered in this guide, refer to the following documentation:

- *HP Array Configuration Utility User Guide*, available on the software CD that is provided with the server, or downloadable from the HP website.
- *HP Servers Troubleshooting Guide*, available on the Documentation CD that is provided with the server.
- *HP ROM-Based Setup Utility User Guide*, available on the Documentation CD that is provided with the server, or downloadable from the HP website.

Getting Help

If you have a problem and have exhausted the information in this guide, you can get further information and other help in the following locations.

Technical Support

In North America, call the HP Technical Support Phone Center at 1-800-652-6672. This service is available 24 hours a day, 7 days a week. For continuous quality improvement, calls may be recorded or monitored. Outside North America, call the nearest HP Technical Support Phone Center. Telephone numbers for worldwide Technical Support Centers are listed on the HP website, <http://www.hp.com>.

Be sure to have the following information available before you call HP:

- Technical support registration number (if applicable)
- Product serial number
- Product model name and number
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software

- Operating system type and revision level

HP Website

The HP website has information on this product as well as the latest drivers and flash ROM images. You can access the HP website at <http://www.hp.com>.

Authorized Reseller

For the name of your nearest authorized reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.
- Elsewhere, see the HP website for locations and telephone numbers.

Reader's Comments

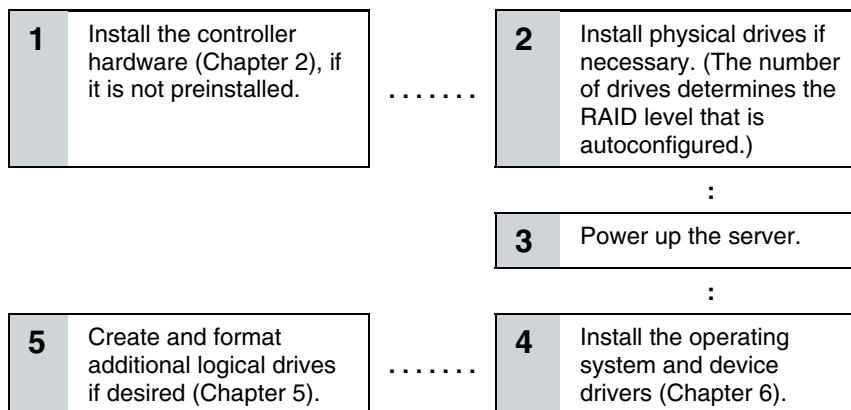
HP welcomes your comments on this guide. Send your comments and suggestions to ServerDocumentation@hp.com.

Installation Overview

The recommended procedure for installing the controller depends on whether the server has been configured, and whether it can self-configure during its first power-up. (To determine whether a server is autoconfigurable, refer to the server-specific setup and installation guide.) The flowcharts on the following pages summarize the recommended procedure for each situation.

Never-Used, Autoconfigurable Server

IMPORTANT: To be sure that the server will autoconfigure the way that you want it to, study the details of the autoconfiguration process in the server-specific setup and installation guide before powering up the server.



Never-Used, Non-Autoconfigurable Server

1	Install the controller hardware (Chapter 2), if it is not preinstalled.
----------	---

:

2	Update the system firmware (Chapter 3).
----------	---

:

3	Update the controller firmware (Chapter 3).
----------	---

:

4	Configure the server (Chapter 4).
----------	-----------------------------------

:

5	Create at least one logical drive and format it (Chapter 5).
----------	--

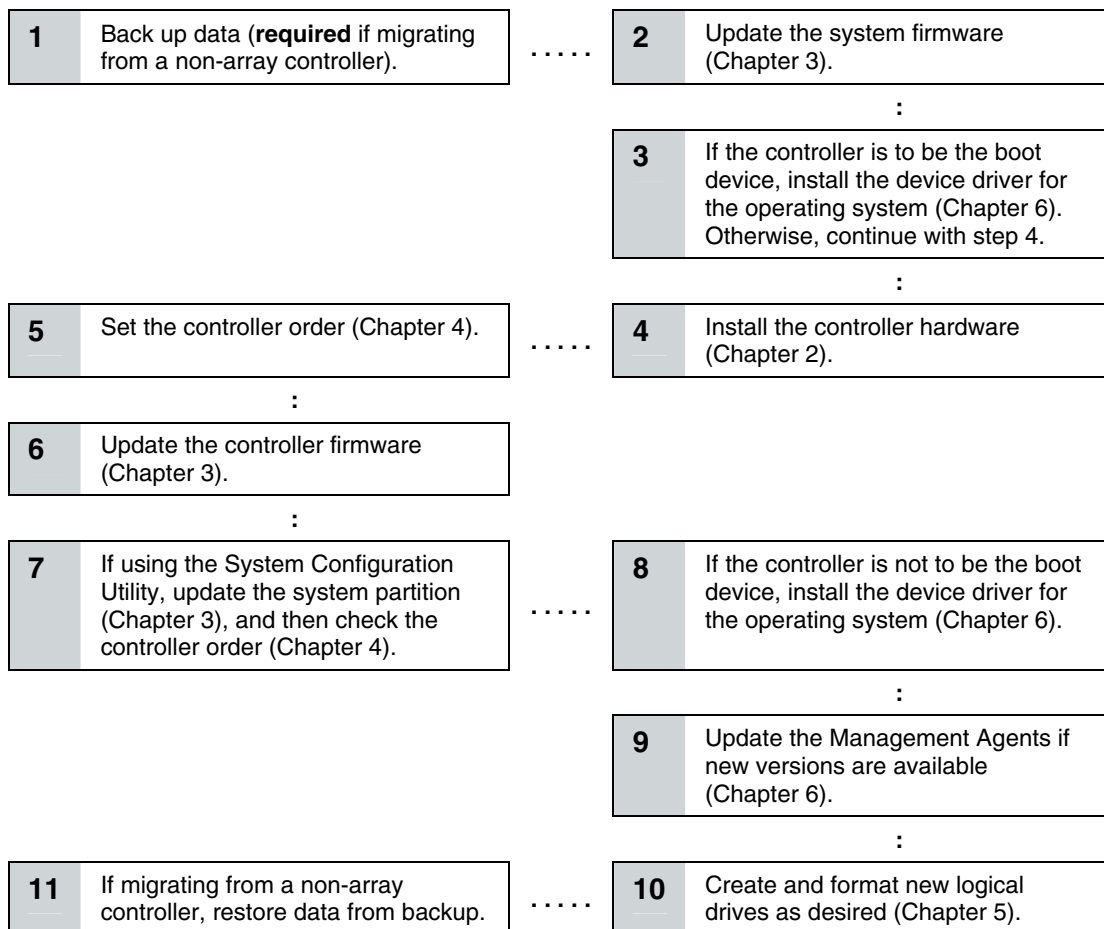
:

6	Install the operating system and device drivers (Chapter 6).
----------	--

:

7	Create and format additional logical drives if desired (Chapter 5).
----------	---

Previously Configured Server



Installing the Hardware

Before beginning the installation procedure, visit the HP website, <http://www.hp.com/support>, to confirm that you have the latest version of each driver and utility file needed. Compare the version numbers of the files there with those of the same files on the software CD or DVD that is supplied in the controller kit.

Preparing the Server

Before installing the controller in the server, back up all data. This step is **required** if you are moving non-arrayed SCSI drives to a Smart Array controller, because data is not preserved during a move between array controllers and non-array controllers.

If the server supports hot-pluggable devices, you can install the controller board without any further preparation (refer to the “Installing the Controller Board” section).

To prepare a server that does not support hot-pluggable devices:

1. Close all applications.
2. Power down the server.



CAUTION: In systems using external data storage, be sure that the server is the first unit powered down and the last unit to be powered back up. This precaution ensures that the system will not erroneously mark the drives as failed.

3. Power down any peripheral devices that are attached to the server.
4. Unplug the AC power cord from the outlet, and then from the server.
5. Disconnect any peripheral devices from the server.

Installing the Controller Board



WARNING: To reduce the risk of personal injury or damage to the equipment, consult the safety information and user documentation provided with the computer before attempting the installation.

Many computers are capable of providing energy levels that are considered hazardous, and are only intended to be serviced by qualified personnel trained in dealing with these hazards. Do not remove enclosures or attempt to bypass any interlocks that may be provided for the purpose of removing these hazardous conditions.

1. Remove or open the access panel.
2. Select an available 3.3-V PCI or PCI-X slot.
3. Remove the slot cover or open the hot-plug latch. Save the retaining screw, if one is present.
4. Slide the controller board along the slot alignment guide, and press the board firmly into the slot so that the contacts on the board edge are properly seated in the system board connector.
5. Secure the controller board in place with the hot-plug latch or retaining screw. If there is a guide latch on the rear of the board, close the latch.
6. To finish installing the hardware, connect the internal and external drives by following the instructions given in “Connecting Storage Devices.”

Connecting Storage Devices

The controller supports Ultra2, Ultra160 (Ultra3), and Ultra320 drives.

Each peripheral that is connected to the controller must have a unique SCSI ID value within the range of 0 to 15 (except ID 7, which is reserved for controller use). This value determines the priority that is given to the device when it attempts to use the SCSI bus.

The system automatically sets the SCSI IDs for hot-pluggable devices if they are supported. For non-hot-pluggable devices, you must set the ID values manually by using switches or jumpers on the device itself.

SCSI buses require termination on both ends to prevent signal degradation. In HP ProLiant servers, however, the controller, SCSI cable, and backplane already provide this termination.

Connecting Internal Storage

1. If the storage device that you are adding is not hot-pluggable, power down the system.
2. Install drives in the removable media bays on the server.



CAUTION: Do not use hot-pluggable drives on the same SCSI bus as non-hot-pluggable drives.

NOTE: Drives that are to be grouped in the same array should have the same capacity.

For additional information about drive installation, refer to the appropriate section in this guide (“Replacing, Moving, or Adding Hard Drives”) and consult the documentation that accompanied the drives.

When you have finished installing drives, continue with the next step.

- If the drives are hot-pluggable, go to step 3.
 - If the drives are not hot-pluggable, go to step 4.
3. Attach the internal point-to-point SCSI cable (provided with the server) from the internal connector of the controller to the hot-plug drive cage.

Installation of the hot-pluggable drives is complete.

4. For each SCSI bus, manually set the SCSI ID on each drive to a unique value in the range of 0 to 15, except 7 (which is reserved for controller use). For detailed instructions, consult the documentation that is provided with the drive.
5. Attach a multi-device SCSI cable from the internal connector of the controller to the non-hot-pluggable hard drives. (The cable may have been provided with the server.)
6. Replace the access panel, and secure it with the thumbscrews if any are present.



CAUTION: Do not operate the server for long periods without the access panel. Operating the server without the access panel results in improper airflow and improper cooling that can lead to thermal damage.

Connecting External Storage (Model 642 only)

1. On the rear of the server, connect the external cable to the VHDCI connector on the controller, and tighten the lock screws on the cable connector.
2. Attach the other end of the cable to the storage enclosure, and tighten the lock screws on the cable connector.

SCSI Cable Part Numbers

Table 2-1: External SCSI Cables for Storage Enclosures

Cable Type	Length	Option Kit Number	Cable Assembly Number
External, VHDCI to VHDCI	1.8 m / 6 ft	341174-B21	313374-001
	3.6 m / 12 ft	341175-B21	313374-002
	7.2 m / 24 ft	164604-B21	313374-004
	11.7 m / 39 ft	150214-B21	313374-005
Internal, multi-device	--	166389-B21	148785-001

Note: If additional cables are required, order by the option kit number.

Updating the Firmware

To update the firmware, you can use the Smart Components (also known as Online ROM Flash Components) that are available on the HP website, <http://www.hp.com/support/proliantstorage>.

1. Locate the Smart Components for the operating system and controller that the server is using.
2. Follow the instructions for installing the components. These instructions are given on the same Web page as the components.
3. Follow the additional instructions that describe how to use the components to flash the ROM. These instructions are provided with each component.

Alternatively, you can use the software CD that is provided in the controller kit. Printed instructions are provided with the CD. Because the Smart Components may be more recent than the firmware upgrade files on the CD, check the Smart Components on the website before using the updates on the CD.

IMPORTANT: If you are updating the firmware on a system that was configured using SCU, you must update the system partition immediately after you have finished updating the firmware. For more information, refer to “Configuring the Server.”

Configuring the Server

After installing the controller hardware and updating the firmware, configure the server by using either RBSU or SCU. For more information, refer to the *HP ROM-Based Setup Utility User Guide* or the server-specific setup and installation guide.

Using RBSU

RBSU is a system configuration utility that is embedded in the system ROM. It is customized for the server on which it is installed.



CAUTION: Not all servers support RBSU. Do not flash an RBSU-ROM image onto a server that is already configured with SCU unless the update instructions specifically state that upgrading from SCU to RBSU is supported. If the upgrade is not supported, the consequences of upgrading are unpredictable, and you may lose data.

1. Power up the server.
2. Press the **F9** key when prompted during system startup.
The main RBSU screen is displayed.
3. Configure the system. (For detailed instructions, refer to the *HP ROM-Based Setup Utility User Guide*.)
4. Select **Boot Controller Order** on the main RBSU screen and follow the on-screen prompts to set the boot controller.
5. When you have finished using RBSU, press the **Esc** key, and then press the **F10** key to confirm that you want to exit. The server reboots with the new configuration.

Using SCU

If you updated the firmware in a used system that was not configured using RBSU, you must use SCU immediately afterward to update the system partition.

1. Locate the page on the HP website (<http://www.hp.com/support>) that contains SCU, and then follow the on-screen instructions to create four SCU diskettes.
2. Insert SCU diskette #1 into the server diskette drive.
3. Restart the system.
4. Select **System Configuration Utility** from the menu or list of icons that is displayed.
5. Follow the on-screen instructions to update or create and populate a system partition.
6. Exit from SCU. If the server does not reboot or a CD error message is displayed, press the **Ctrl+Alt+Del** keys to reboot the server manually.

When you have finished using SCU to configure the system, use ORCA immediately afterward to confirm that the controller order is unchanged, as follows:

1. Reboot the server. The POST sequence begins, and an ORCA prompt message is briefly displayed.
2. Press the **F8** key to start ORCA.

NOTE: The ORCA prompt is displayed for only a few seconds. If you do not press the **F8** key during this time, you must restart the server to obtain the prompt again.

3. On the **Main Menu** screen, select **Select as Boot Controller**.
4. Follow the remaining prompts to set the currently selected controller as the boot controller for the system.

If you want to use ORCA to create logical drives, you do not need to exit the utility at this point. Continue as described in Chapter 5.

Configuring an Array

HP provides two utilities for manually configuring an array on a Smart Array controller:

- **Array Configuration Utility (ACU)**—A versatile, browser-based utility that provides maximum control over configuration parameters
- **Option ROM Configuration for Arrays (ORCA)**—A simple ROM-based configuration utility that runs on all operating systems

NOTE: To copy a particular array configuration to several other servers on the same network, use the Array Configuration Replicator (ACR) or the scripting capability of ACU. ACR is provided in the SmartStart Scripting Toolkit, available at <http://www.hp.com/servers/sstoolkit>.

Whichever utility you use, the following limitations apply:

- For the most efficient use of drive space, do not mix drives of different capacities within the same array. The configuration utility treats all physical drives in an array as if they have the same capacity as the smallest drive in the array. The excess capacity of any larger drives is wasted because it is unavailable for data storage.
- The probability that an array will experience a hard drive failure increases with the number of hard drives in the array. If you configure a logical drive with RAID 5, keep the probability of failure low by using no more than 14 physical drives in the array.

Table 5-1: Comparison of Utilities for Configuring an Array

	ACU*	ORCA*
Uses a graphical interface	y	n
Available in languages other than English	y	n
Executable at any time	y	n
Available on CD	y	n
Suggests the optimum configuration for an unconfigured controller	y	n
Describes configuration errors	y	n
Supports these operating systems:		
Microsoft® Windows® Server 2003	y	y
Microsoft Windows 2000	y	y
Microsoft Windows NT®	y	y
Novell NetWare	y [†]	y
Linux	y	y
Allows these procedures:		
Creation and deletion of arrays, logical drives	y	y
Assignment of RAID level	y	y
Sharing of spare drive among several arrays	y	n
Assignment of multiple spare drives per array	y	n
Setting of stripe size	y	n
Migration of RAID level or stripe size	y	n
Configuration of controller settings	y	n
Expansion of an array	y	n
Creation of multiple logical drives per array	y	n
Setting of boot controller	n	y
*y = feature is supported; n = feature is not supported		
[†] NetWare can use ACU only when the system is offline.		

For conceptual information about arrays, logical drives, and fault-tolerance methods, refer to Appendix D.

Using ACU

For detailed information about using ACU, refer to the *HP Array Configuration Utility User Guide*. This document is available on the Documentation CD that is provided in the controller kit.

Using ORCA

When a server is powered up, the Power-On Self-Test (POST) runs, and any array controllers that are in the system are initialized. If the array controller supports ORCA, POST temporarily halts, and an ORCA prompt message is displayed for approximately five seconds. (If ORCA is not supported, the prompt message is not displayed, and the system continues with the startup sequence.)

While the prompt is displayed, press the **F8** key to start ORCA. The ORCA main menu is displayed, allowing you to create, view, or delete a logical drive. (On a ProLiant system, you can also use ORCA to set the currently selected controller as the boot controller.)




Figure 5-1: ORCA main menu screen

Configuration Procedure

To create a logical drive using ORCA:

1. Select **Create Logical Drive**.

The screen displays a list of all available (unconfigured) physical drives and the valid RAID options for the system.

2. Use the **Arrow** keys, **Spacebar**, and **Tab** key to navigate around the screen and set up the logical drive, including an online spare drive if one is required.

NOTE: You cannot use ORCA to configure one spare drive to be shared among several arrays. Only ACU enables you to configure shared spare drives.

While configuring the logical drive, one of the settings allows you to use either 4 GB or 8 GB as the maximum boot drive size. Selecting 8 GB allows a larger boot partition for operating systems such as Microsoft Windows NT 4.0 that use cylinders, heads, and sectors of a physical drive to determine the drive size. The larger boot drive size also lets you increase the size of the logical drive at some later time. However, logical drive performance is likely to decrease if the larger boot drive size is enabled.

3. Press the **Enter** key to accept the settings.
4. Press the **F8** key to confirm the settings and save the new configuration.

After several seconds, the **Configuration Saved** screen is displayed.

5. Press the **Enter** key to continue.

You can now create another logical drive by repeating the previous steps.

NOTE: Newly created logical drives are invisible to the operating system. To make the new logical drives available for data storage, format them using the instructions given in the operating system documentation.

Installing the Device Drivers and Management Agents

Device Drivers

The drivers for the controller are located on the Support Software CD or the SmartStart CD that is provided in the controller kit. Updates are posted to the HP website, <http://www.hp.com/support>.

Using the Support Software CD: Instructions for installing the drivers from the Support Software CD are given in the leaflet that is supplied with the CD. Note that the exact procedure depends on whether the server is new or already contains the operating system and user data.

Using the SmartStart CD: If you use the Assisted Installation path feature of SmartStart to install the operating system on a new server, the drivers are automatically installed at the same time.

You can also use SmartStart to update the drivers manually on older systems. For further information, refer to the SmartStart documentation.

Management Agents

If you use the Assisted Installation path feature of SmartStart to install the operating system on a new server, the Management Agents are automatically installed at the same time.

You can update the Management Agents on older servers by using the latest versions of the agents from one of these sources:

- The Management CD, obtainable from your local HP reseller or authorized service provider
- The SmartStart CD
- The HP website, <http://www.hp.com/servers/manage>

For the procedure to update the agents, refer to the documentation on the Management CD or on the HP website.

If the new agents do not function correctly, you may also need to update Insight Manager. The latest versions of Insight Manager are also available for download at the HP website.


Upgrading or Replacing the Cache




WARNING: There is a risk of explosion, fire, or personal injury if the battery pack in the cache is not properly handled. Refer to the Battery Replacement Notice in Appendix A before installing or removing the cache.

To install a cache module:

1. Pull out the latches on the memory socket (1).
2. Insert the cache module straight into the socket (2) applying equal pressure on both ends of the module until the latches snap into place. If the latches do not snap into place, press them inward to secure the module in the socket.



3. With the lip on the retainer clip facing the cache module, insert one prong of the clip into the corresponding hole in the controller board just above the cache module (1).
4. Squeeze the prongs together slightly (2) and insert the other prong into the remaining hole (3).



To remove the cache module, reverse this procedure.

Regulatory Compliance Notices

Federal Communications Commission Notice

Part 15 of the Federal Communications Commission (FCC) Rules and Regulations has established Radio Frequency (RF) emission limits to provide an interference-free radio frequency spectrum. Many electronic devices, including computers, generate RF energy incidental to their intended function and are, therefore, covered by these rules. These rules place computers and related peripheral devices into two classes, A and B, depending upon their intended installation. Class A devices are those that may reasonably be expected to be installed in a business or commercial environment. Class B devices are those that may reasonably be expected to be installed in a residential environment (for example, personal computers). The FCC requires devices in both classes to bear a label indicating the interference potential of the device as well as additional operating instructions for the user.

The rating label on the device shows the classification (A or B) of the equipment. Class B devices have an FCC logo or FCC ID on the label. Class A devices do not have an FCC logo or FCC ID on the label. After the Class of the device is determined, refer to the corresponding statement in the following sections.

Class A Equipment

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

Class B Equipment

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit that is different from that to which the receiver is connected.
- Consult the dealer or an experienced radio or television technician for help.

Declaration of Conformity for Products Marked with the FCC Logo, United States Only

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

For questions regarding your product, contact us by mail or telephone:

- Hewlett-Packard Company
P. O. Box 692000, Mail Stop 530113
Houston, Texas 77269-2000
- 1-800-652-6672 (For continuous quality improvement, calls may be recorded or monitored.)

For questions regarding this FCC declaration, contact us by mail or telephone:

- Hewlett-Packard Company
P. O. Box 692000, Mail Stop 510101
Houston, Texas 77269-2000
- 1-281-514-3333

To identify this product, refer to the part, series, or model number found on the product.

Modifications

The FCC requires the user to be notified that any changes or modifications made to this device that are not expressly approved by Hewlett-Packard Company may void the user's authority to operate the equipment.

Cables

Connections to this device must be made with shielded cables with metallic RFI/EMI connector hoods in order to maintain compliance with FCC Rules and Regulations.

Canadian Notice (Avis Canadien)

Class A Equipment

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Class B Equipment

This Class B digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

European Union Notice

Products with the CE Marking comply with both the EMC Directive (89/336/EEC) and the Low Voltage Directive (73/23/EEC) issued by the Commission of the European Community.

Compliance with these directives implies conformity to the following European Norms (the equivalent international standards are in parentheses):

- EN55022 (CISPR 22) – Electromagnetic Interference
- EN55024 (IEC61000-4-2, 3, 4, 5, 6, 8, 11) – Electromagnetic Immunity
- EN60950 (IEC950) – Product Safety

BSMI Notice

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Japanese Notice

ご使用になっている装置にVCCIマークが付いていましたら、次の説明文をお読み下さい。

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスB情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。

VCCIマークが付いていない場合には、次の点にご注意下さい。

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean Notices

A급 기기 (업무용 정보통신기기)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약 잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

B급 기기 (가정용 정보통신기기)

이 기기는 가정용으로 전자파적합등록을 한 기기로서
주거지역에서는 물론 모든지역에서 사용할 수 있습니다.

Battery Replacement Notice

The battery-backed write cache uses a nickel metal hydride (NiMH) battery pack.



WARNING: There is a risk of explosion, fire, or personal injury if the battery pack is not properly handled. To reduce this risk:

- Do not try to recharge the batteries if they are disconnected from the controller.
- Do not expose the battery pack to water, or to temperatures higher than 60°C.
- Do not abuse, disassemble, crush, or puncture the battery pack.
- Do not short the external contacts.
- Replace the battery pack only with the designated HP spare.

Battery disposal should comply with local regulations. Alternatively, use established parts return methods to return the battery pack to HP for disposal.



Batteries, battery packs, and accumulators should not be disposed of together with the general household waste. To forward them to recycling or proper disposal, use the public collection system or return them to HP, your authorized HP Partners, or their agents.

For more information about battery replacement or proper disposal, contact your HP authorized reseller or your authorized service provider.

Electrostatic Discharge

To prevent damaging the system, be aware of the precautions you need to follow when setting up the system or handling parts. A discharge of static electricity from a finger or other conductor may damage system boards or other static-sensitive devices. This type of damage may reduce the life expectancy of the device.

To prevent electrostatic damage, observe the following precautions:

- Avoid hand contact by transporting and storing products in static-safe containers.
- Keep electrostatic-sensitive parts in their containers until they arrive at static-free workstations.
- Place parts on a grounded surface before removing them from their containers.
- Avoid touching pins, leads, or circuitry.
- Always be properly grounded when touching a static-sensitive component or assembly.

There are several methods for grounding. Use one or more of the following methods when handling or installing electrostatic-sensitive parts:

- Use a wrist strap connected by a ground cord to a grounded workstation or computer chassis. Wrist straps are flexible straps with a minimum of 1 megohm resistance in the ground cords. To provide proper ground, wear the strap snug against the skin.
- Use heel straps, toe straps, or boot straps at standing workstations. Wear the straps on both feet when standing on conductive floors or dissipating floor mats.
- Use conductive field service tools.

- Use a portable field service kit with a folding static-dissipating work mat.

If you do not have any of the suggested equipment for proper grounding, have an HP authorized reseller install the part.

NOTE: For more information on static electricity, or assistance with product installation, contact your HP authorized reseller.

Controller Specifications

Table C-1: Controller Specifications

Dimensions (PCB only)	28.7 cm × 10.8 cm × 1.6 cm (11.3 in × 4.3 in × 0.63 in)
Power required	No more than 14 W for either model
PCI-X bus transfer rate	Up to 1064 MB/s at 133 MHz (64-bit)
Temperature range	Operating: 10° to 35°C (50° to 95°F) Storage: –30° to 60°C (–22° to 140°F)
Relative humidity (noncondensing)	Operating: 20% to 80% Storage: 5% to 90%
RAID levels supported	0, 1+0, 5
SCSI Bus:	
Number of channels	1 internal (641); 1 internal and 1 external (642)
Maximum number of drives per channel	15 (14 in an external enclosure)
Connector type	68-pin Wide internal, VHDCI external
Termination	Required, but provided on Compaq and newer HP systems
Transfer rate	Up to 320 MB/s (80 MHz) per channel

For more information about the controller features and specifications, refer to <http://www.compaq.com/smartarray>.

Drive Arrays and Fault Tolerance

What Is a Drive Array?

The capacity and performance of a single physical (hard) drive is adequate for home users. However, business users demand higher storage capacities, higher data transfer rates, and greater protection against data loss when drives fail.

Connecting extra physical drives (P_n in the figure) to a system increases the total storage capacity, but has no effect on the efficiency of read/write (R/W) operations. Data can still be transferred to only one physical drive at a time.




Figure D-1: Physical drives added to system

With an array controller installed in the system, the capacity of several physical drives can be combined into one or more virtual units called **logical drives** (also called logical volumes and denoted by L_n in the figures in this section). Then, the read/write heads of all the constituent physical drives are active simultaneously, reducing the total time required for data transfer.




Figure D-2: Physical drives configured into a logical drive (L1)

Because the read/write heads are active simultaneously, the same amount of data is written to each drive during any given time interval. Each unit of data is called a **block** (denoted by B_n in Figure D-3), and adjacent blocks form a set of data **stripes** (S_n) across all the physical drives that comprise the logical drive.




Figure D-3: Data striping (S1–S4) of data blocks B1–B12

For data in the logical drive to be readable, the data block sequence must be the same in every stripe. This sequencing process is performed by the array controller, which sends the data blocks to the drive write heads in the correct order.

A natural consequence of the striping process is that each physical drive in a given logical drive contains the same amount of data. If one physical drive has a larger capacity than other physical drives in the same logical drive, the extra capacity is wasted because it cannot be used by the logical drive.

The group of physical drives containing the logical drive is called a **drive array**, or just **array** (denoted by A_n in Figure D-4). Since all the physical drives in an array are commonly configured into just one logical drive, the term array is also often used as a synonym for logical drive. However, an array can contain several logical drives, each of a different size.

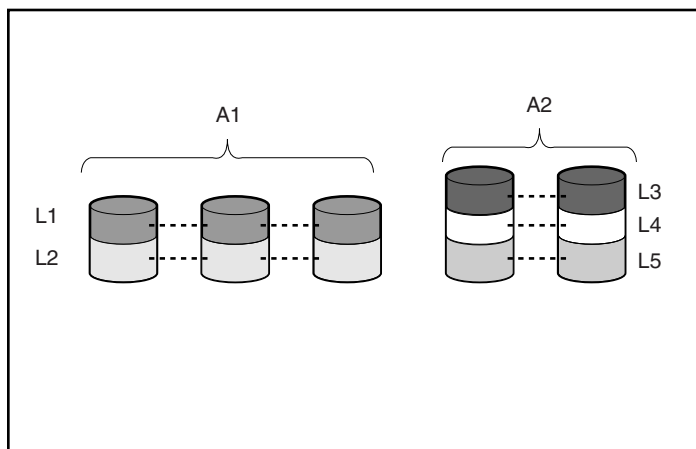


Figure D-4: Two arrays (A1, A2) containing five logical drives spread across five physical drives

Each logical drive in an array is distributed across all of the physical drives within the array. A logical drive can also extend across more than one port on the same controller, but it cannot extend across more than one controller.

Drive failure, although rare, is potentially catastrophic. In Figure D-4, for example, failure of **any** physical drive causes **all** logical drives in the same array to fail, and all data on the drives is lost. To protect against data loss due to physical drive failure, logical drives are configured with **fault tolerance**. For more information, refer to “Fault-Tolerance Methods.”

For any configuration except RAID 0, further protection against data loss can be achieved by assigning a drive as an **online spare** (or **hot spare**). This drive contains no data and is connected to the same controller as the array. When any other physical drive in the array fails, the controller automatically rebuilds information that was originally on the failed drive to the online spare. The system is thus restored to full RAID-level data protection, although it now no longer has an online spare. (However, in the unlikely event that another drive in the array fails while data is being rewritten to the spare, the logical drive will still fail.)

When you configure an online spare, it is automatically assigned to all logical drives in the same array. Additionally, you do not need to assign a separate online spare to each array. Instead, you can configure one hard drive to be the online spare for several arrays if the arrays are all on the same controller.

Fault-Tolerance Methods

Several fault-tolerance methods exist. Those that are most often used with Smart Array controllers are hardware-based RAID methods.

Two alternative fault-tolerance methods that are sometimes used are described in the section “Alternative Fault-Tolerance Methods.” However, hardware-based RAID methods provide a much more robust and controlled fault-tolerance environment, so these alternative methods are seldom used.

Hardware-Based Fault-Tolerance Methods

The hardware-based methods that are recommended for use with Smart Array controllers are:

- RAID 0—Data Striping only (no fault tolerance)
- RAID 1+0—Drive Mirroring
- RAID 5—Distributed Data Guarding
- RAID ADG—Advanced Data Guarding

RAID 0—No Fault Tolerance

A RAID 0 configuration (refer to Figure D-3 for an example) provides data striping, but there is no protection against data loss when a drive fails. However, it is useful for rapid storage of large amounts of non-critical data (for printing or image editing, for example), or when cost is the most important consideration.

Advantages

- Has the highest write performance of all RAID methods.
- Has the lowest cost per unit of stored data of all RAID methods.
- All drive capacity is used to store data (none is needed for fault tolerance).

Disadvantages

- All data on the logical drive is lost if a physical drive fails.
- Cannot use an online spare.
- Can only preserve data by backing it up to external drives.

RAID 1+0—Drive Mirroring

In a RAID 1+0 configuration, data is duplicated to a second drive.




Figure D-5: Drive mirroring of P1 to P2

When the array has more than two physical drives, drives are mirrored in pairs.




Figure D-6: Mirroring with more than two physical drives in the array

In each mirrored pair, the physical drive that is not busy answering other requests answers any read requests that are sent to the array. (This behavior is called **load balancing**.) If a physical drive fails, the remaining drive in the mirrored pair can still provide all the necessary data. Several drives in the array can fail without incurring data loss, as long as no two failed drives belong to the same mirrored pair.

This fault-tolerance method is useful when high performance and data protection are more important than the cost of physical drives.

NOTE: When there are only two physical drives in the array, this fault-tolerance method is often referred to as RAID 1.

Advantages

- Has the highest read and write performance of any fault-tolerant configuration.
- No data is lost as long as no failed drive is mirrored to another failed drive (up to half of the physical drives in the array can fail).

Disadvantages

- This method is expensive (many drives are needed for fault tolerance).
- Only half of the total drive capacity is usable for data storage.

RAID 5—Distributed Data Guarding

In a RAID 5 configuration, data protection is provided by **parity data** (denoted by $P_{x,y}$ in Figure D-7). This parity data is calculated stripe by stripe from the user data that is written to all other blocks within that stripe. The blocks of parity data are distributed evenly over every physical drive within the logical drive.




Figure D-7: Distributed data guarding, showing parity information ($P_{x,y}$)

When a physical drive fails, data that was on the failed drive can be calculated from the remaining parity data and user data on the other drives in the array. This recovered data is usually written to an online spare in a process called a **rebuild**.

This configuration is useful when cost, performance, and data availability are equally important.

Advantages

- Has high read performance.
- Data is not lost if only one physical drive fails.
- More drive capacity is usable than with RAID 1+0—parity information requires only the storage space equivalent to one physical drive.

Disadvantages

- Has relatively low write performance.
- Data is lost if a second drive fails before data from the first failed drive is rebuilt.

RAID ADG—Advanced Data Guarding

NOTE: Not all controllers support RAID ADG.

RAID ADG, like RAID 5, generates and stores parity information to protect against data loss caused by drive failure. With RAID ADG, however, two different sets of parity data are used (denoted by $P_{x,y}$ and $Q_{x,y}$ in Figure D-8), allowing data to still be preserved if two drives fail. Each set of parity data uses a capacity equivalent to that of one of the constituent drives.




Figure D-8: Advanced data guarding (RAID ADG)

This method is most useful when data loss is unacceptable, but cost is also an important factor. The probability that data loss will occur when an array is configured with RAID ADG is less than it would be if it were configured with RAID 5.

Advantages

- Has high read performance.
- Allows high data availability—any two drives can fail without loss of critical data.
- More drive capacity is usable than with RAID 1+0—parity information requires only the storage space equivalent to two physical drives.

Disadvantage

The main disadvantage of RAID ADG is a relatively low write performance (lower than RAID 5), because of the need for two sets of parity data.

Comparison of RAID Methods

Table D-1 summarizes the important features of the different kinds of RAID methods described here. The decision chart in Table D-2 may help you to determine which option is best for your situation.

Table D-1: Summary of RAID Methods

	RAID 0	RAID 1+0	RAID 5	RAID ADG*
Alternative name	Striping (no fault tolerance)	Mirroring	Distributed Data Guarding	Advanced Data Guarding
Usable drive space**	100%	50%	67% to 93%	50% to 96%
Usable drive space formula	n	$n/2$	$(n-1)/n$	$(n-2)/n$
Minimum number of physical drives	1	2	3	4
Tolerates failure of one physical drive?	No	Yes	Yes	Yes
Tolerates simultaneous failure of more than one physical drive?	No	Only if no two failed drives are in the same mirrored pair	No	Yes

continued

Table D-1: Summary of RAID Methods *continued*

	RAID 0	RAID 1+0	RAID 5	RAID ADG*
Read performance	High	High	High	High
Write performance	High	Medium	Low	Low
Relative cost	Low	High	Medium	Medium

*Not all controllers support RAID ADG.

**Values for usable drive space are calculated with these assumptions: (1) all physical drives in the array have the same capacity; (2) online spares are not used; (3) no more than 14 physical drives are used per array for RAID 5; (4) no more than 56 drives are used with RAID ADG.

Table D-2: Choosing a RAID Method

Most Important	Also Important	Suggested RAID Level
Fault tolerance	Cost effectiveness	RAID ADG*
	I/O performance	RAID 1+0
Cost effectiveness	Fault tolerance	RAID ADG*
	I/O performance	RAID 5 (RAID 0 if fault tolerance is not required)
I/O performance	Cost effectiveness	RAID 5 (RAID 0 if fault tolerance is not required)
	Fault tolerance	RAID 1+0

*Not all controllers support RAID ADG.

Alternative Fault-Tolerance Methods

Your operating system may also support software-based RAID or controller duplexing.

- **Software-based RAID** resembles hardware-based RAID, except that the operating system works with logical drives as if they were physical drives. To protect against data loss caused by physical drive failure, each logical drive must be in a different array from the others.

- **Controller Duplexing** uses two identical controllers with independent, identical sets of drives containing identical data. In the unlikely event of a controller failure, the remaining controller and drives will service all requests.

Neither of these alternative fault-tolerance methods supports online spares or automatic data recovery, nor do they support auto-reliability monitoring or interim data recovery.

If you decide to use one of these alternative fault-tolerance methods, configure your arrays with RAID 0 for maximum storage capacity and refer to your operating system documentation for further implementation details.

Replacing, Moving, or Adding Hard Drives

Identifying the Status of a Hard Drive

When a drive is configured as part of an array and attached to a powered-up controller, the condition of the drive can be determined from the illumination pattern of the hard drive status lights (LEDs). Table E-1 describes the meanings of the different illumination patterns in a ProLiant server.




Figure E-1: Hard drive status LED indicators

Table E-1: Hard Drive Status from LED Illumination Pattern

Activity LED (1)	Online LED (2)	Fault LED (3)	Interpretation
On, off, or flashing	On or off	Flashing	A predictive failure alert has been received for this drive. Replace the drive as soon as possible.
On, off, or flashing	On	Off	The drive is online and configured as part of an array. If the array is configured for fault tolerance and all other drives in the array are online, and a predictive failure alert is received or a drive capacity upgrade is in progress, you may replace the drive online.
On or flashing	Flashing	Off	Do not remove the drive. Removing a drive may terminate the current operation and cause data loss. The drive is rebuilding or undergoing capacity expansion.
On	Off	Off	Do not remove the drive. The drive is being accessed, but (1) it is not configured as part of an array; (2) it is a replacement drive and rebuild has not yet started; or (3) it is spinning up during the Power-On Self-Test (POST) sequence.
Flashing	Flashing	Flashing	Do not remove the drive. Removing a drive may cause data loss in non-fault-tolerant configurations. Either (1) the drive is part of an array being selected by an array configuration utility; (2) Drive Identification has been selected in Insight Manager; or (3) drive firmware is being updated.
Off	Off	On	The drive has failed and has been placed offline. You may replace the drive.
Off	Off	Off	Either (1) the drive is not configured as part of an array; (2) the drive is configured as part of an array, but it is a replacement drive that is not being accessed or being rebuilt yet; or (3) the drive is configured as an online spare. If the drive is connected to an array controller, you may replace the drive online.

Recognizing Hard Drive Failure

In a ProLiant system, a steadily glowing Fault LED indicates that that drive has failed. Other means by which hard drive failure is revealed are:

- The amber LED on the front of a storage system illuminates if failed drives are inside. (However, this LED also illuminates when other problems occur, such as when a fan fails, a redundant power supply fails, or the system overheats.)
- A POST message lists failed drives whenever the system is restarted, as long as the controller detects at least one functional drive.
- ACU represents failed drives with a distinctive icon.
- Insight Manager can detect failed drives remotely across a network. (For more information about Insight Manager, refer to the documentation on the Management CD.)
- The Array Diagnostic Utility (ADU) lists all failed drives.

For additional information about diagnosing hard drive problems, refer to the *HP Servers Troubleshooting Guide*.



CAUTION: Sometimes, a drive that has previously been failed by the controller may seem to be operational after the system is power-cycled or (for a hot-pluggable drive) after the drive has been removed and reinserted. However, continued use of such marginal drives may eventually result in data loss. Replace the marginal drive as soon as possible.

Effects of a Hard Drive Failure

When a hard drive fails, all logical drives that are in the same array are affected. Each logical drive in an array may be using a different fault-tolerance method, so each logical drive can be affected differently.

- RAID 0 configurations cannot tolerate drive failure. If any physical drive in the array fails, all non-fault-tolerant (RAID 0) logical drives in the same array will also fail.
- RAID 1+0 configurations can tolerate multiple drive failures as long as no failed drives are mirrored to one another.

- RAID 5 configurations can tolerate one drive failure.
- RAID ADG configurations can tolerate simultaneous failure of two drives.

Compromised Fault Tolerance

If more hard drives fail than the fault-tolerance method allows, fault tolerance is compromised, and the logical drive fails. In this case, all requests from the operating system are rejected with unrecoverable errors. You are likely to lose data, although it can sometimes be recovered (refer to “Recovering from Compromised Fault Tolerance”).

One example of a situation in which compromised fault tolerance may occur is when a drive in an array fails while another drive in the array is being rebuilt. If the array has no online spare, any logical drives in this array that are configured with RAID 5 fault tolerance will fail.

Compromised fault tolerance can also be caused by non-drive problems, such as a faulty cable or temporary power loss to a storage system. In such cases, you do not need to replace the physical drives. However, you may still have lost data, especially if the system was busy at the time that the problem occurred.

Recovering from Compromised Fault Tolerance

If fault tolerance is compromised, inserting replacement drives does not improve the condition of the logical volume. Instead, if the screen displays unrecoverable error messages, perform the following procedure to recover data:

1. Power down the entire system, and then power it back up. In some cases, a marginal drive will work again for long enough to allow you to make copies of important files.

If a 1779 POST message is displayed, press the **F2** key to re-enable the logical volumes. Remember that data loss has probably occurred and any data on the logical volume is suspect.

2. Make copies of important data, if possible.
3. Replace any failed drives.

After you have replaced the failed drives, the fault tolerance may again be compromised. If so, cycle the power again. If the 1779 POST message is displayed:

- a. Press the **F2** key to re-enable the logical drives.
- b. Recreate the partitions.
- c. Restore all data from backup.

To minimize the risk of data loss that is caused by compromised fault tolerance, make frequent backups of all logical volumes.

Replacing Hard Drives

The most common reason for replacing a hard drive is that it has failed. However, another reason is to gradually increase the storage capacity of the entire system.

If you insert a hot-pluggable drive into a drive bay while the system power is on, all disk activity in the array pauses while the new drive is spinning up. This spin-up process usually lasts for approximately 20 seconds. When the drive has achieved its normal spin rate, data recovery to the replacement drive begins automatically (as indicated by the blinking Online LED on the replacement drive) if the array is in a fault-tolerant configuration.

If you replace a drive belonging to a fault-tolerant configuration while the system power is off, a POST message is displayed when the system is next powered up. This message prompts you to press the **F1** key to start automatic data recovery. If you do not enable automatic data recovery, the logical volume remains in a ready-to-recover condition and the same POST message is displayed whenever the system is restarted.

Factors to Consider Before Replacing Hard Drives

- In systems that use external data storage, be sure that the server is the first unit to be powered down and the last to be powered back up. Taking this precaution ensures that the system does not erroneously mark the drives as failed when the server is powered up.

- If you set the SCSI ID jumpers manually:
 - Check the ID value of the removed drive to be sure that it corresponds to the ID of the drive marked as failed.
 - Set the same ID value on the replacement drive to prevent SCSI ID conflicts.

Before replacing a degraded drive:

- Open Insight Manager and inspect the Error Counter window for each physical drive in the same array to confirm that no other drives have any errors. (For details, refer to the Insight Manager documentation on the Management CD.)
- Be sure that the array has a current, valid backup.
- Use replacement drives that have a capacity at least as great as that of the smallest drive in the array. The controller immediately fails drives that have insufficient capacity.

To minimize the likelihood of fatal system errors, take these precautions when removing failed drives:

- Do not remove a degraded drive if any other drive in the array is offline (the Online LED is off). In this situation, no other drive in the array can be removed without data loss.

These cases are the exceptions:

- When RAID 1+0 is used, drives are mirrored in pairs. Several drives can be in a failed condition simultaneously (and they can all be replaced simultaneously) without data loss, as long as no two failed drives belong to the same mirrored pair.
- When RAID ADG is used, two drives can fail simultaneously (and be replaced simultaneously) without data loss.
- If the offline drive is a spare, the degraded drive can be replaced.
- Do not remove a second drive from an array until the first failed or missing drive has been replaced **and** the rebuild process is complete. (The rebuild is complete when the Online LED on the front of the drive stops blinking.)

These cases are the exceptions:

- In RAID ADG configurations, any two drives in the array can be replaced simultaneously.
- In RAID 1+0 configurations, any drives that are not mirrored to other removed or failed drives can be simultaneously replaced offline without data loss.

Automatic Data Recovery (Rebuild)

When you replace a hard drive in an array, the controller uses the fault-tolerance information on the remaining drives in the array to reconstruct the missing data (the data that was originally on the replaced drive) and write it to the replacement drive. This process is called automatic data recovery, or rebuild. If fault tolerance is compromised, this data cannot be reconstructed and is likely to be permanently lost.

If another drive in the array fails while fault tolerance is unavailable during rebuild, a fatal system error may occur, and all data on the array is then lost. In exceptional cases, however, failure of another drive need not lead to a fatal system error. These exceptions include:

- Failure after activation of a spare drive
- Failure of a drive that is not mirrored to any other failed drives (in a RAID 1+0 configuration)
- Failure of a second drive in a RAID ADG configuration

Time Required for a Rebuild

The time required for a rebuild varies considerably, depending on several factors:

- The priority that the rebuild is given over normal I/O operations (you can change the priority setting by using ACU)
- The amount of I/O activity during the rebuild operation
- The rotational speed of the hard drives
- The availability of drive cache

- The brand, model, and age of the drives
- The amount of unused capacity on the drives
- The number of drives in the array (for RAID 5 and RAID ADG)

Allow approximately 15 minutes per gigabyte for the rebuild process to be completed. This figure is conservative, and newer drive models usually require less time to rebuild.

System performance is affected during the rebuild, and the system is unprotected against further drive failure until the rebuild has finished. Therefore, replace drives during periods of low activity when possible.



CAUTION: If the Online LED of the replacement drive stops blinking and the amber Fault LED glows, or if other drive LEDs in the array go out, the replacement drive has failed and is producing unrecoverable disk errors. Remove and replace the failed replacement drive.

When automatic data recovery has finished, the Online LED of the replacement drive stops blinking and begins to glow steadily.

Failure of Another Drive During Rebuild

If a non-correctable read error occurs on another physical drive in the array during the rebuild process, the Online LED of the replacement drive stops blinking and the rebuild abnormally terminates.

If this situation occurs, reboot the server. The system may temporarily become operational for long enough to allow recovery of unsaved data. In any case, locate the faulty drive, replace it, and restore data from backup.

Upgrading Hard Drive Capacity

You can increase the storage capacity on a system even if there are no available drive bays by swapping drives one at a time for higher capacity drives. This method is viable as long as a fault-tolerance method is running.



CAUTION: Because it can take up to 15 minutes per gigabyte to rebuild the data in the new configuration, the system is unprotected against drive failure for many hours while a given drive is upgraded. Perform drive capacity upgrades only during periods of minimal system activity.

To upgrade hard drive capacity:

1. Back up all data.
2. Replace any drive. The data on the new drive is recreated from redundant information on the remaining drives.



CAUTION: Do not replace any other drive until data rebuild on this drive is complete.

3. When data on the new drive has been rebuilt (the Activity LED turns off), repeat the previous step for the other drives in the array, one at a time.

When you have replaced all drives, you can use the extra capacity to either create new logical drives or extend existing logical drives. For more information about these procedures, refer to the *HP Array Configuration Utility User Guide*.

Moving Drives and Arrays

You can move drives to other ID positions on the same array controller. You can also move a complete array from one controller to another, even if the controllers are on different servers.

Before you move drives, the following conditions must be met:

- The move will not result in more than 14 physical drives per controller channel.
- No controller will be configured with more than 32 logical volumes.
- The array has no failed or missing drives.
- The array is in its original configuration.
- The controller is not reading from, or writing to, any of the spare drives in the array.

- The controller is not running capacity expansion, capacity extension, or RAID or stripe size migration.
- The controller is using the latest firmware version (recommended).

If you want to move an array to another controller, you must also consider the following additional limitations:

- All drives in the array must be moved at the same time.
- In most cases, a moved array (and the logical drives that it contains) can still undergo array capacity expansion, logical drive capacity extension, or migration of RAID level or stripe size. An exception occurs when the array meets **all** of these conditions:
 - It was originally created on a SMART-2/P, SMART-2DH, SA-3200, SA-3100ES, SA-4200, SA-4250ES, or SA-530x controller.
 - It is moved to a controller that does not have a battery-backed cache.
 - It has less than 4 MB of unused capacity.
- If a controller contains a RAID ADG logical volume, none of the arrays on the controller can be moved **directly** to a controller that does not support RAID ADG. (The arrays can be moved **indirectly**, as described by the instructions in this section.)

When all the conditions have been met:

1. Back up all data before removing any drives or changing configuration. This step is **required** if you are moving data-containing drives from a controller that does not have a battery-backed cache.
2. Power down the system.
3. If you are moving an array from a controller that contains a RAID ADG logical volume to a controller that does not support RAID ADG:
 - a. Remove or disconnect the drives that contain the RAID ADG logical volume.
 - b. Reboot the server.
 - c. Open ACU and navigate to the controller that contained the RAID ADG volume.

ACU displays the missing RAID ADG volume using a different icon to indicate that the volume is unavailable.

- d. Delete the RAID ADG volume.
 - e. Accept the configuration change, and then close ACU.
 - f. Power down the system.
4. Move the drives.
 5. Power up the system. If a 1724 POST message is displayed, drive positions were changed successfully and the configuration was updated.
If a 1785 POST message is displayed:
 - a. Power down the system immediately to prevent data loss.
 - b. Return the drives to their original locations.
 - c. Restore the data from backup, if necessary.
 6. Check the new drive configuration by running ORCA or ACU.

Adding Drives

You can add hard drives to a system at any time, as long as you do not exceed the maximum number of drives that the controller supports. You can then either build a new array from the added drives or use the extra storage capacity to expand the capacity of an existing array.

To perform an array capacity expansion, use ACU. If the system is using hot-pluggable drives, you can expand array capacity without shutting down the operating system (that is, with the server online) if ACU is running in the same environment as the normal server applications. (For more information, refer to the *HP Array Configuration Utility User Guide*.)

The expansion process is illustrated in the following figure, in which the original array (containing data) is shown with a dashed border and the newly added drives (containing no data) are shown unshaded. The array controller adds the new drives to the array and redistributes the original logical drives over the enlarged array one logical drive at a time. This process liberates some storage capacity on each of the physical drives in the array. During this procedure, the logical drives each keep the same fault-tolerance method in the enlarged array that they had in the smaller array.




Figure E-2: Array capacity expansion

When the expansion process has finished, you can use the liberated storage capacity on the enlarged array to create new logical drives. Alternatively, you can enlarge one of the original logical drives. This latter process is called **logical drive capacity extension** and is also carried out using ACU.

Probability of Logical Drive Failure

The probability that a logical drive will fail depends on the RAID level setting and on the number and type of physical drives in the array. If the logical drive does not have an online spare, the following results apply.

- A RAID 0 logical drive fails if only one physical drive fails.
- A RAID 1+0 logical drive fails if any two failed physical drives are mirrored to each other.
 - The **maximum** number of physical drives that can fail **without** causing failure of the logical drive is $n/2$, where n is the number of hard drives in the array. In practice, a logical drive usually fails before this maximum is reached. As the number of failed physical drives increases, it becomes increasingly likely that the newly failed drive is mirrored to a previously failed drive.
 - The **minimum** number of physical drive failures that can cause the logical drive to fail is two. This situation occurs when the two failed drives are mirrored to each other. As the total number of drives in the array increases, the probability that the only two failed drives in an array are mirrored to each other decreases.
- A RAID 5 logical drive fails if two physical drives fail.
- A RAID ADG logical drive fails when three physical drives fail.

At any given RAID level, the probability of logical drive failure increases as the number of physical drives in the logical drive increases. This is illustrated more quantitatively in Figure F-1. The data for this graph is calculated from the mean time between failures (MTBF) value for a typical physical drive, assuming that no online spares are present. If an online spare is added to any of the fault-tolerant RAID configurations, the probability of logical drive failure is further decreased.




Figure F-1: Relative probability of logical drive failure

Troubleshooting

Several diagnostic tools provide feedback about problems with arrays. The most important are:

- ADU

This utility can be downloaded from the HP website, <http://www.hp.com/support>. The meanings of the various ADU error messages are provided in the *HP Servers Troubleshooting Guide*.

- POST Messages

Smart Array controllers produce diagnostic error messages at reboot. Many of these POST messages are self-explanatory and suggest corrective actions. For more information about POST messages, refer to the *HP Servers Troubleshooting Guide*.

- Server Diagnostics

To use Server Diagnostics:

- a. Insert the SmartStart CD into the server CD-ROM drive.
- b. Click **Agree** when the license agreement is displayed, and select the **Maintenance** tab.
- c. Click **Server Diagnostics**, and follow the on-screen prompts and instructions.

Index

A

- ACR (Array Configuration Replicator) 5-1
- ACU (Array Configuration Utility) 5-1
- ADG *See* RAID ADG
- ADU (Array Diagnostics Utility) G-1
- advanced data guarding *See* RAID ADG
- alert, predictive failure E-2
- array
 - adding hard drives to E-11
 - defined D-3
 - manual configuration of, using ORCA 5-3
 - mixing drive capacities in 5-1
 - moving E-9
 - online spares in D-5
 - physical limitations of D-4
- array capacity expansion E-11
- Array Configuration Replicator (ACR) 5-1
- Array Configuration Utility (ACU) 5-1
- array controller
 - configuration of 5-1
 - dimensions of C-1
 - driver installation for 6-1
 - duplexing of D-12
 - installation of 1-1, 2-1
 - power requirements of C-1
- Array Diagnostics Utility (ADU) G-1
- authorized reseller x

- automatic data recovery
 - description of E-7
 - limitation of D-12

B

- batteries, recycling A-6
- block of data, defined D-2
- boot controller, setting 4-1
- boot straps, using B-1

C

- cables
 - FCC compliance statement for A-3
 - part numbers for 2-4
- cabling instructions 2-2
- capacity extension E-12
- capacity upgrade of hard drives E-8
- comparison
 - of ACU with ORCA 5-2
 - of different RAID methods D-10
 - of hardware-based RAID with software-based RAID D-11
 - of logical drive failure risk for different RAID levels F-3
 - of RAID methods with other fault-tolerance methods D-11

- configuring
 - array controller 5-1
 - SCSI ID settings 2-2
 - server 4-1
- controller
 - configuration of 5-1
 - dimensions of C-1
 - driver installation for 6-1
 - duplexing of D-12
 - installation of 1-1, 2-1
 - power requirement of C-1
- controller installation
 - flowcharts for 1-1
 - precautions during 2-2
- controller order, setting 4-1

D

- data block, defined D-2
- data protection methods
 - non-RAID D-11
 - RAID D-5
- data rebuild time E-7
- data recovery, automatic E-7
- data stripes, defined D-2
- data transfer rate C-1
- Declaration of Conformity A-3
- device drivers, installing 6-1
- device priority, setting 2-2
- diagnosing problems
 - error messages in POST G-1
 - general G-1
 - hard drive E-3
- dimensions of controller C-1
- distributed data guarding (RAID 5) D-8
- drive array *See* array
- drive failure
 - POST notification of E-3
 - probability graph F-3
 - replacing drive after E-5
- drive mirroring (RAID 1+0) D-6
- drive status LEDs E-1
- drivers, installing and updating 6-1

- drives of different capacity on array 5-1
- duplexing, controller D-12

E

- ESD (electrostatic discharge) B-1
- expanding an array E-11
- extending a logical drive E-12
- external drives 2-4
- external storage, powering up and down 2-1

F

- fault tolerance *See also* RAID methods
 - alternative methods of D-11
 - compromised E-4
 - controller duplexing as D-12
 - description of methods D-5
 - software-based RAID as D-11
- FCC notices A-1
- features
 - of ACU 5-2
 - of controller C-1
 - of ORCA 5-2
 - of RAID methods D-10
- Federal Communications Commission
 - notices A-1
- firmware, updating 3-1
- flowcharts, controller installation 1-1

G

- graph, drive failure probability F-3
- grounding methods B-1

H

- hard drive failure
 - detection of E-3
 - fault tolerance and D-10
 - multiple, simultaneous D-10
 - protection against D-5
 - recognizing E-1
 - replacing drive after E-5

hard drive status LEDs, interpreting pattern of E-2

hard drives

adding, to array E-11

capacity of, restrictions on 5-1

different capacity of, on array 5-1

failure of E-3

interpreting status LEDs on E-2

larger capacity, using, in array E-8

LEDs of E-1

minimum number of, for RAID D-10

moving E-9

replacing E-5

status lights on E-1

upgrading capacity of E-8

heel straps, using B-1

hot spare D-4

HP website x

I

Insight Manager 6-1

installing

controller hardware 2-1

controller, flowcharts for 1-1

device drivers 6-1

interim data recovery, limitation of D-12

internal drives, connecting 2-3

J

jumpers, setting 2-2

L

LEDs on hard drives E-1

load balancing, defined D-7

logical drive capacity extension E-12

logical drives

compared to array D-3

creation of, with ORCA 5-3

defined D-2

enlarging (extending) E-12

failure of E-4

recovery of, options for E-4

M

manual configuration of array 5-3

maximum number of hard drives

for RAID 5 D-10

for RAID ADG D-10

minimum number of hard drives for

RAID D-10

mirroring of drives D-6

moving drives E-9

multiple hard drive failure D-10

N

no fault tolerance (RAID 0) D-5

O

online drive capacity upgrade E-8

online spare D-4

option kit part numbers for cables 2-4

Option ROM Configuration for Arrays *See* ORCA

options ROM, updating 3-1

ORCA (Option ROM Configuration for Arrays) 5-1

overview of installation process 1-1

P

parity data

in RAID 5 D-8

in RAID ADG D-9

part numbers for cables 2-4

parts, handling and storing B-1

peripherals, SCSI ID of 2-2

- physical drives *See* hard drives
- POST messages G-1
- power requirements of controller C-1
- powering system up and down, caution for 2-1
- precautions
 - against ESD B-1
 - for installing controller 2-2
 - for setting SCSI IDs 2-2
- predictive failure alert E-2
- protecting data
 - alternative methods D-11
 - RAID methods D-5

R

- RAID 0 (no fault tolerance) D-5
- RAID 1+0 (drive mirroring) D-6
- RAID 5 (distributed data guarding) D-8
- RAID ADG (advanced data guarding) D-9
- RAID methods *See also* fault tolerance
 - comparison with alternative fault-tolerance methods D-11
 - comparison with each other D-10
 - selection chart for D-11
 - software-based D-11
 - summary of features D-10
- RBSU (ROM-Based Setup Utility) 4-1
- rebuild
 - description of E-7
 - time required for E-7
- recovering data, general information for E-4
- replacing hard drive E-5
- reprogramming ROM 3-1
- resources
 - ACR (Array Configuration Replicator) 5-1
 - ACU 5-1
 - ADU (Array Diagnostics Utility) G-1
 - automatic data recovery E-7
 - Insight Manager 6-1
 - Management Agents 6-1

- ORCA 5-1
- POST G-1
- RBSU 4-1
- ROM, updating 3-1
- ROM-Based Setup Utility (RBSU) 4-1

S

- SCSI bus termination requirement 2-3
- SCSI bus, termination of C-1
- SCSI IDs, setting 2-2
- server configuration 4-1
- SmartStart CD, updating firmware using 3-1
- spare drives, defined D-4
- static-safe containers B-1
- status LEDs, on drives E-1
- striping data, defined D-2
- summary of RAID method features D-10
- Support Software CD, updating firmware using 3-1
- symbols in text viii
- system ROM, updating 3-1

T

- technical support ix
- telephone numbers ix, x
- termination of SCSI bus 2-3, C-1
- time needed for data rebuild E-7
- troubleshooting *See also* POST messages
 - general G-1
 - hard drive problems E-1

U

- updating
 - device drivers 6-1
 - firmware 3-1
 - Management Agents 6-1
- upgrading hard drive capacity E-8

utilities

ACR (Array Configuration
Replicator) 5-1

ACU 5-1

Array Diagnostics Utility G-1

Insight Manager 6-1

ORCA 5-1

POST G-1

RBSU 4-1

W

website, HP x

wrist straps B-1